HyperTransport™ in Networking Applications

Van Lewing
Business Development Manager, PMC-Sierra

PMC-SIERRA
Member of HyperTransport Consortium

HyperTransport™
CONSORTIUM

High-Speed I/O in CPU-Based Network Applications

- Traditional PC or server app:
  - Read dataset
  - Crunch, crunch, crunch (in cache)
  - Update dataset
- Network application:
  - Read packet
  - Do various table lookups
  - Update packet
  - Update statistics
- More sensitive to memory & I/O performance

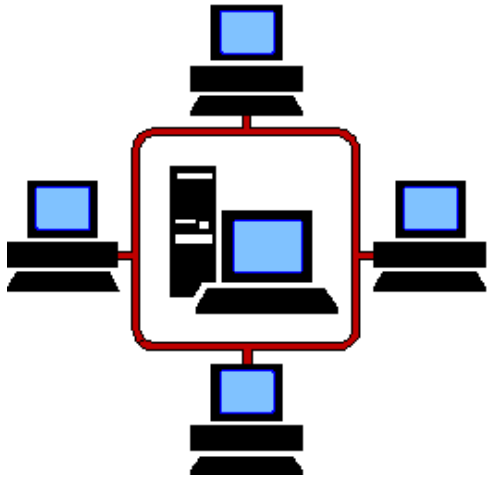

HyperTransport™ Features

- High-performance, scalable, multi-drop interconnect
- Chained and branching topologies
  - Up to 15 secondary devices per host interface
  - Each link consists of a pair of unidirectional data paths
  - Individual links connected via repeaters or bridges
- Scalable bandwidth
  - 2, 4, 8, 16 or 32 data signals per data path
  - 400 to 1600 Gbit/s symbol rate
  - 100 to 6400 Mbyte/s aggregate bandwidth in each direction
  - Double Data Rate (DDR) clocking
- Enhanced Low Voltage Differential Signaling (LVDS)
  - 1.2V driver supply voltage
  - Supports trace lengths of up to 24" (for 800 Mbit/s operation)
  - Can be used over connectors and short cables

HyperTransport™ Architecture

- Physical Layer
  - Defines physical and electrical characteristics
- Data Link Layer
  - Defines initialization, configuration, data transfer and error control procedures
- Protocol Layer
  - Defines commands, virtual channels and ordering rules
- Transaction Layer
  - Defines actions including reads, writes, synchronization, etc.
- Session Layer
  - Defines power management, system management and interrupt functions

HyperTransport™ Data Path

- Each link consists of a pair of unidirectional data paths
- Each data path consists of:
  - 2, 4, 8, 16 or 32 command/address/data lines (CAD)
  - 1 control line (CTL)
  - 1, 2 or 4 clock lines (CLK)
- Extendable to multiple processors
  - Protocol supports I/O stream and endpoint identification (terminate vs. pass-through)

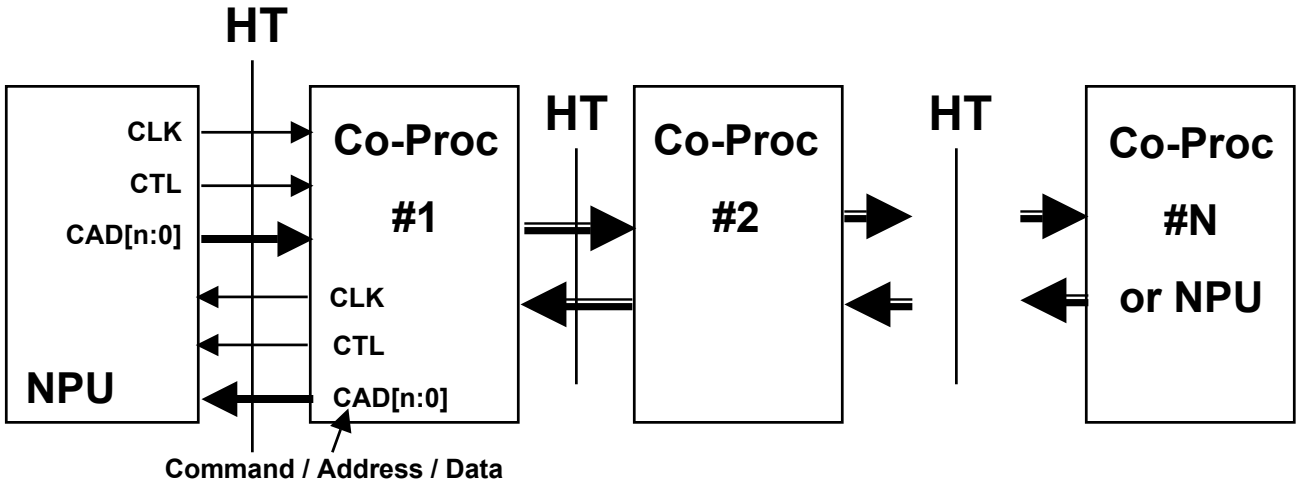

HT Applications

- Chip-to-Chip Interface:
  - CPU to I/O devices (encryption, packet classification, etc.)
  - CPU-to-CPU (connect two SoCs back-to-back)
  - Switched applications (multiple CPUs and I/O devices)

HyperTransport™ Network Usage

- Look-Aside applications
  - HyperTransport already supports a rich set of memory and co-processor access semantics
  - All HyperTransport layers except the Session Layer are directly applicable to Look-Aside applications
- Control Plane applications
  - A common Control Plane interface does not require a separate hardware interface

Control Plane/Data Plane Tasks

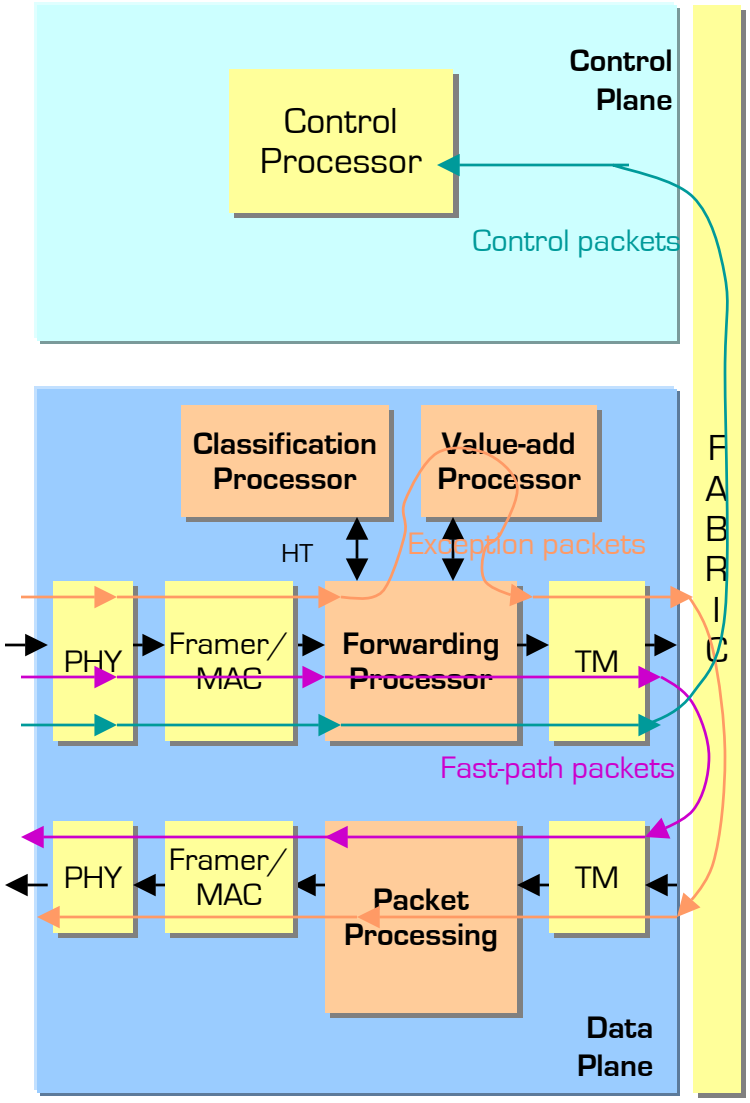

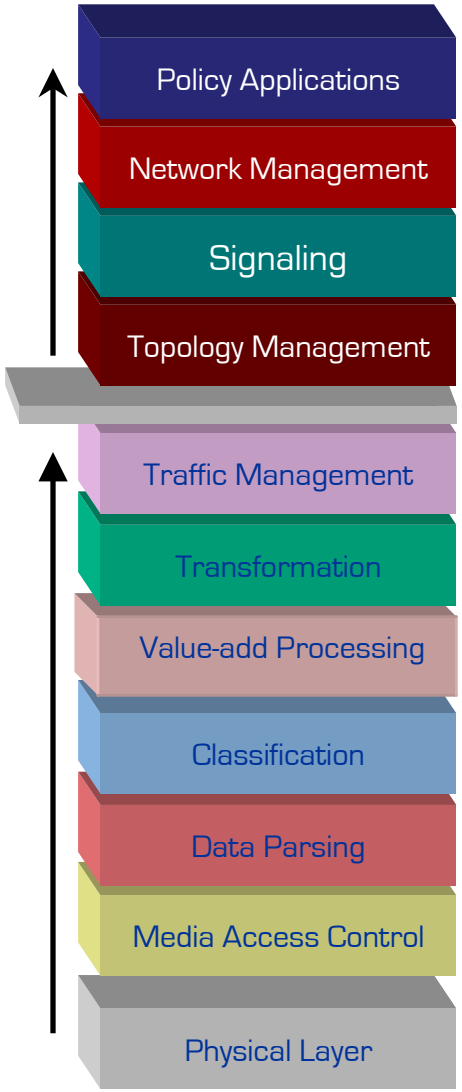

Control Plane tasks: Consists of control and device management functions which are less time critical

Data Plane tasks:
Consists of core functions in the forwarding path of a device

These are time and performance critical

Media access control: Implementation of lower layer protocols, such as Ethernet, SONET framing, ATM cell processing etc

Data parsing: Parsing cell or packet headers, protocol information

Classification: Identifying a packet or cell against a set of criteria defined at L2 or higher of the ISO model

Value-add Processing: Processing of packets with special handling (e.g. tunneling, etc.)

Data transformation: Modification or translation of data within a protocol

Traffic management: Queing, policing, scheduling of data traffic according to QoS, classification and policies

RM9000x2™ Data Bandwidth

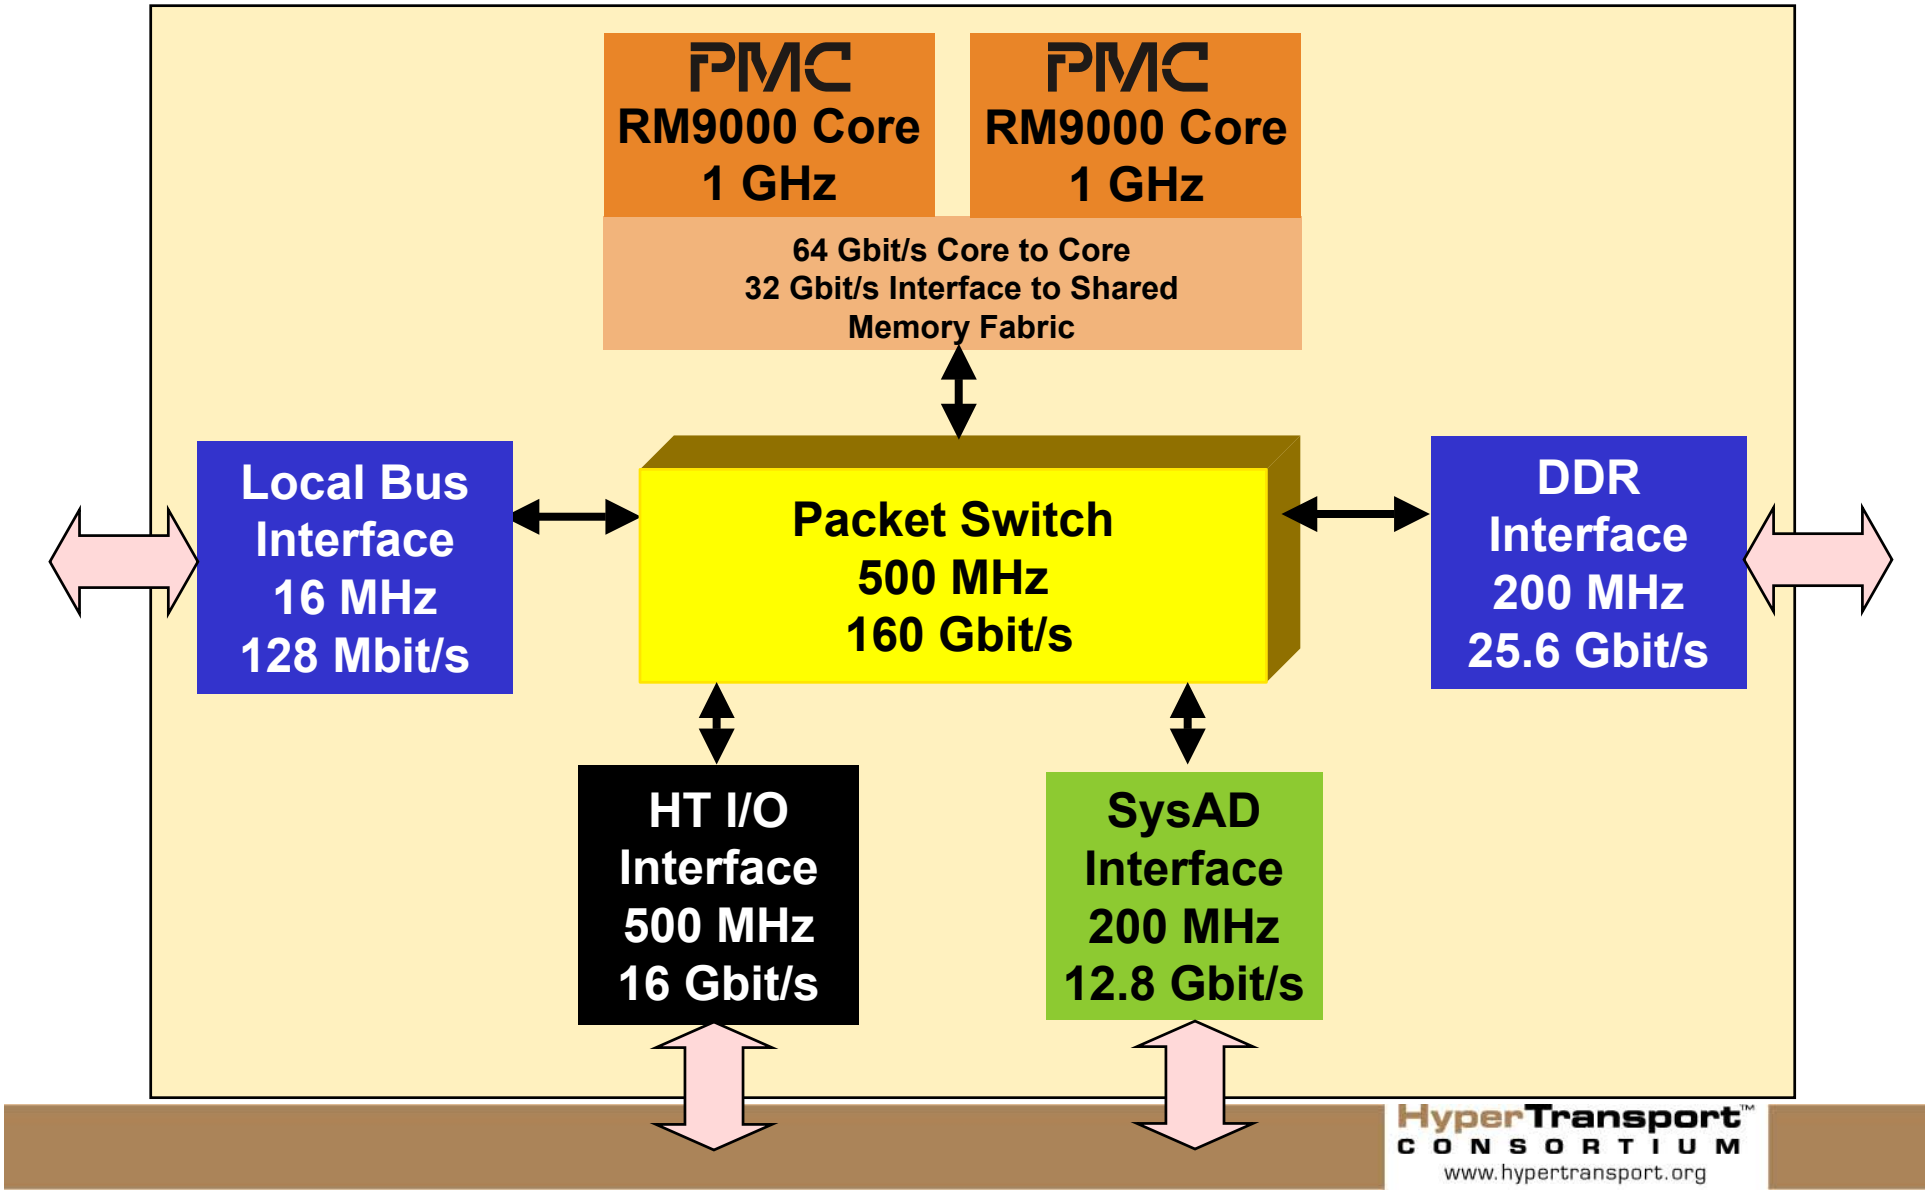

Top-level Software Architecture

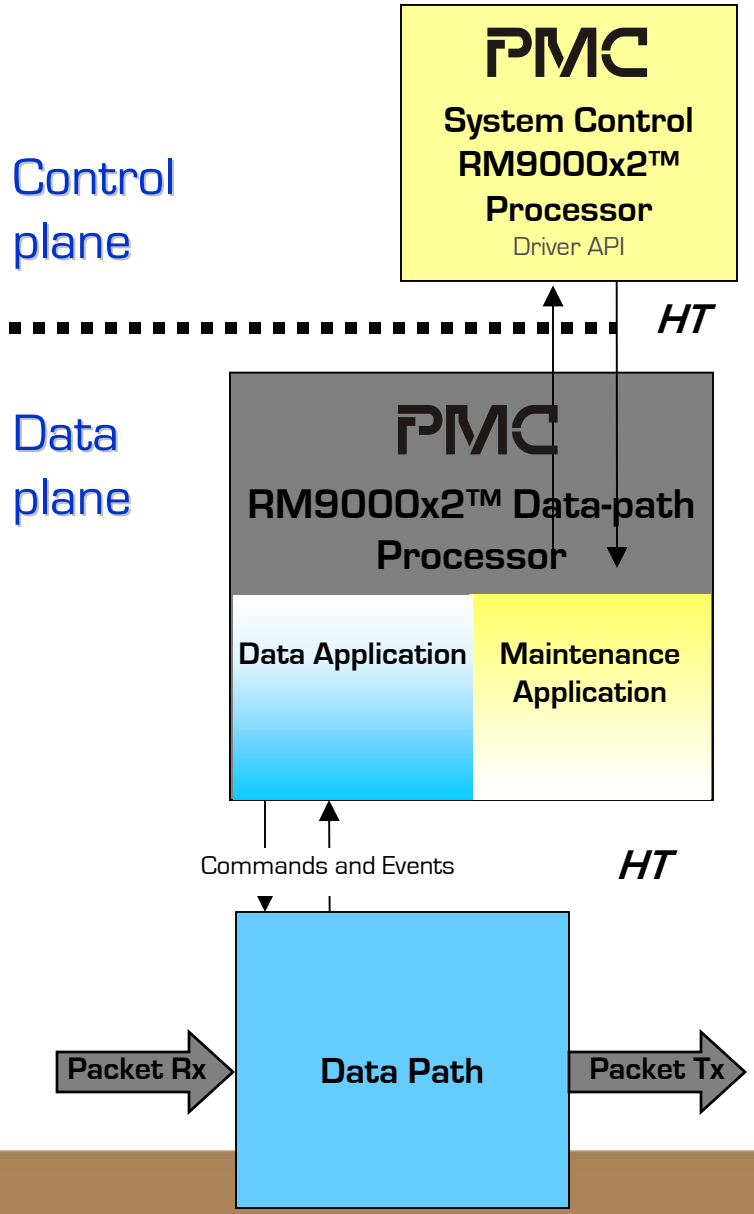

- Datapath CPU communicates with host processor via HT
- PMC-Sierra's RM9000x2™ MIPS-based processor assumed for data-path CPU
- One RM9000x2™ core is dedicated to data path, and other is dedicated entirely to maintenance
- Software on Control plane
  - Device drivers/APIs
  - Data and Maintenance Application firmware
    - MPLS Forwarding
    - FR/ATM Inter-working
    - IPv6
- Communication is with datapath processor only

Streaming vs. Addressable I/O Interfaces

- **Streaming** interfaces read or write an entire packet at a time
  - Simple, but…
    - No preemption once a long packet starts
    - Memory address controlled by DMA engine near the memory

Streaming vs. **Addressable** I/O Interfaces

- **Addressable** interfaces provide a memory address for each transfer
  - Transfer size as small as a single byte
  - Largest individual transfer size is a "block" (typically 128-256 bytes)
  - "Intelligent" peripherals have fine-grained control
  - Peripherals can share data structures with CPUs

Scalability

- Link width and speed can be selected independently by device's designers
  - Wide & fast for high performance
  - Slower & narrower for lower cost
- All HyperTransport™ devices interoperable
- Multiple devices can be switched or cascaded

Bandwidth

- OC-192 Requirements
  - 125 M lookups /sec
  - Key + 4 byte overhead = 22 Gbit/s
  - Key + 8 = 26 Gbit/s
- HyperTransport™ Provides
  - 16-bit @ 1.6 Gbit/s = 25.6 Gbit/s
  - 32-bit @ 800 Mbit/s = 25.6 Gbit/s
  - 32-bit @ 1.6 Gbit/s = 51.2 Gbit/s

HT Increases Bandwidth/Pin

- Point-to-point, unidirectional wiring
  - Separate connections for transmit and receive
  - Simplified electrical termination
  - Eliminate dead states for bus turn-around
- Differential signaling
  - Two wires per logical signal
  - High noise immunity
  - Better tolerance to varying logic thresholds
- Much higher-frequency operation
- Initial HT implementations at 800 Mbit/s per pin pair, targeting over 1.2 Gbit/s

Network CPU possibilities with HyperTransport™

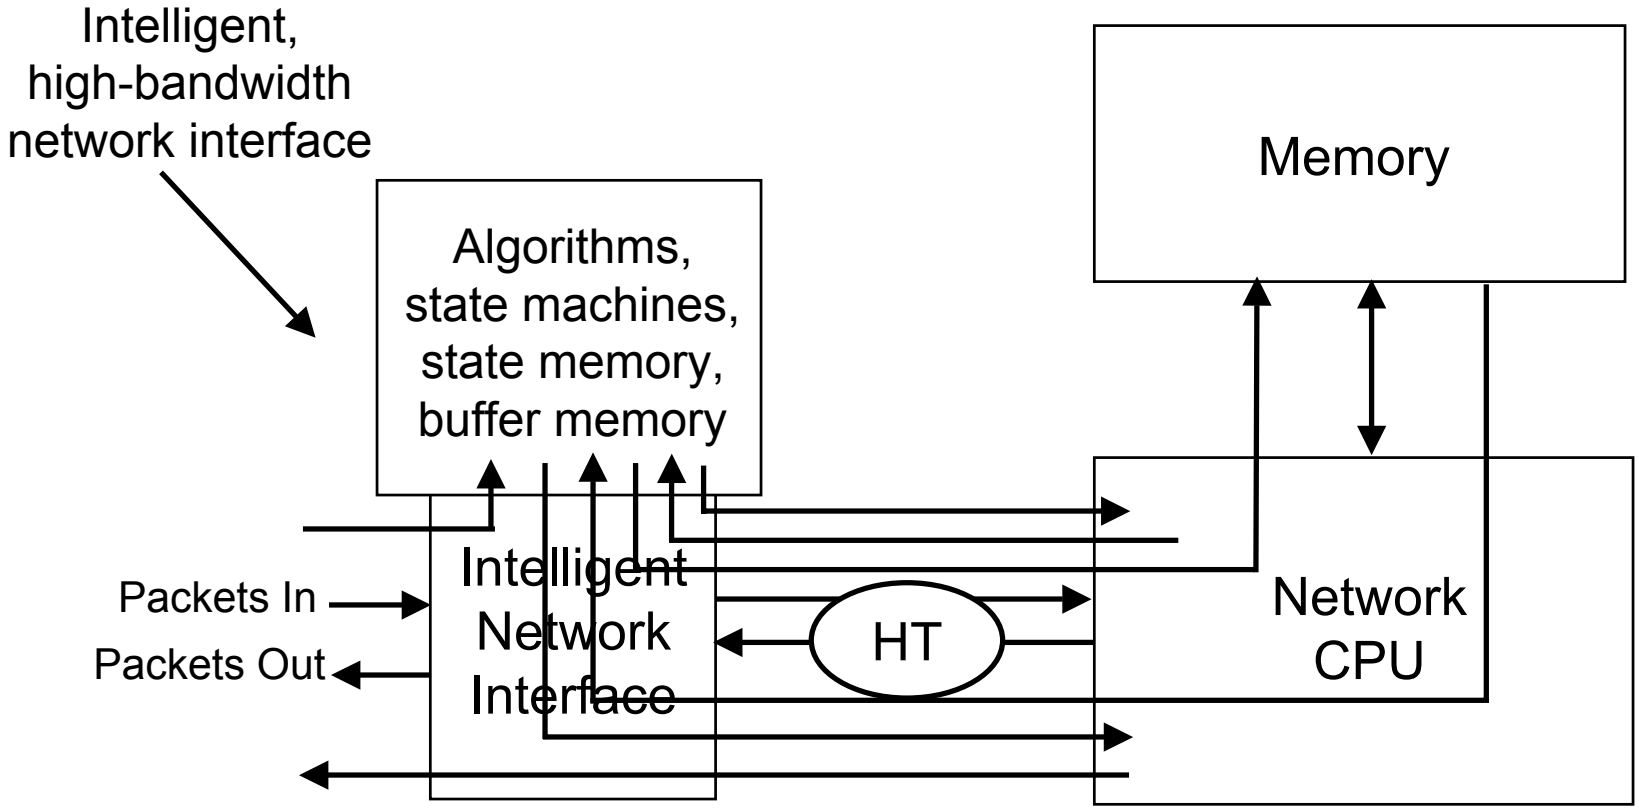

Simple Networking Application

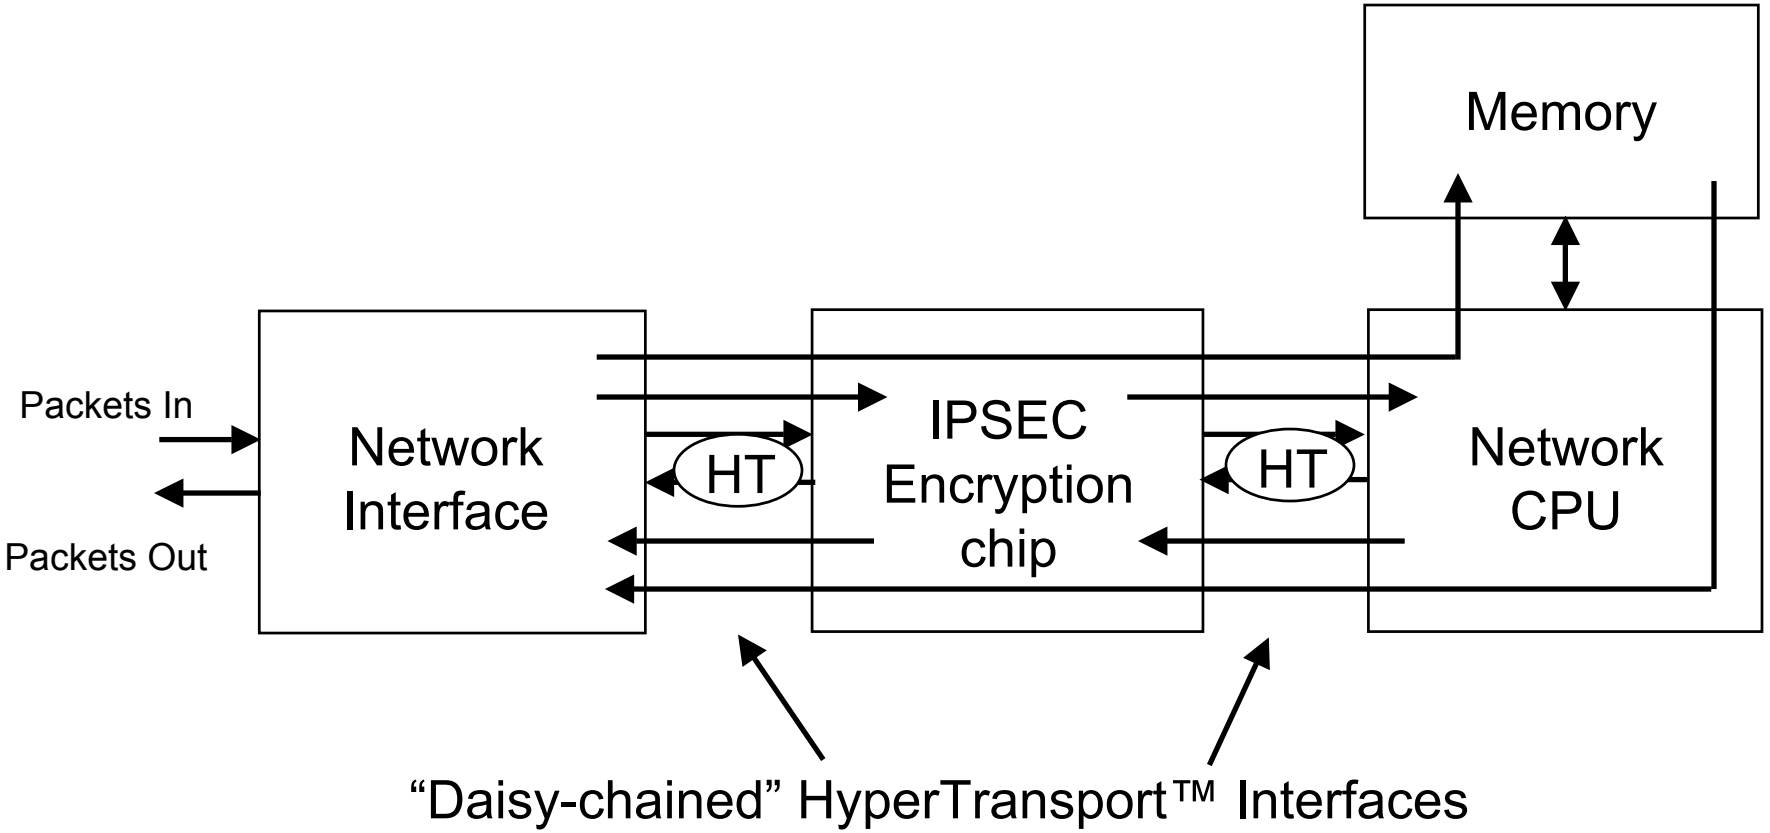

Switched Networking Application

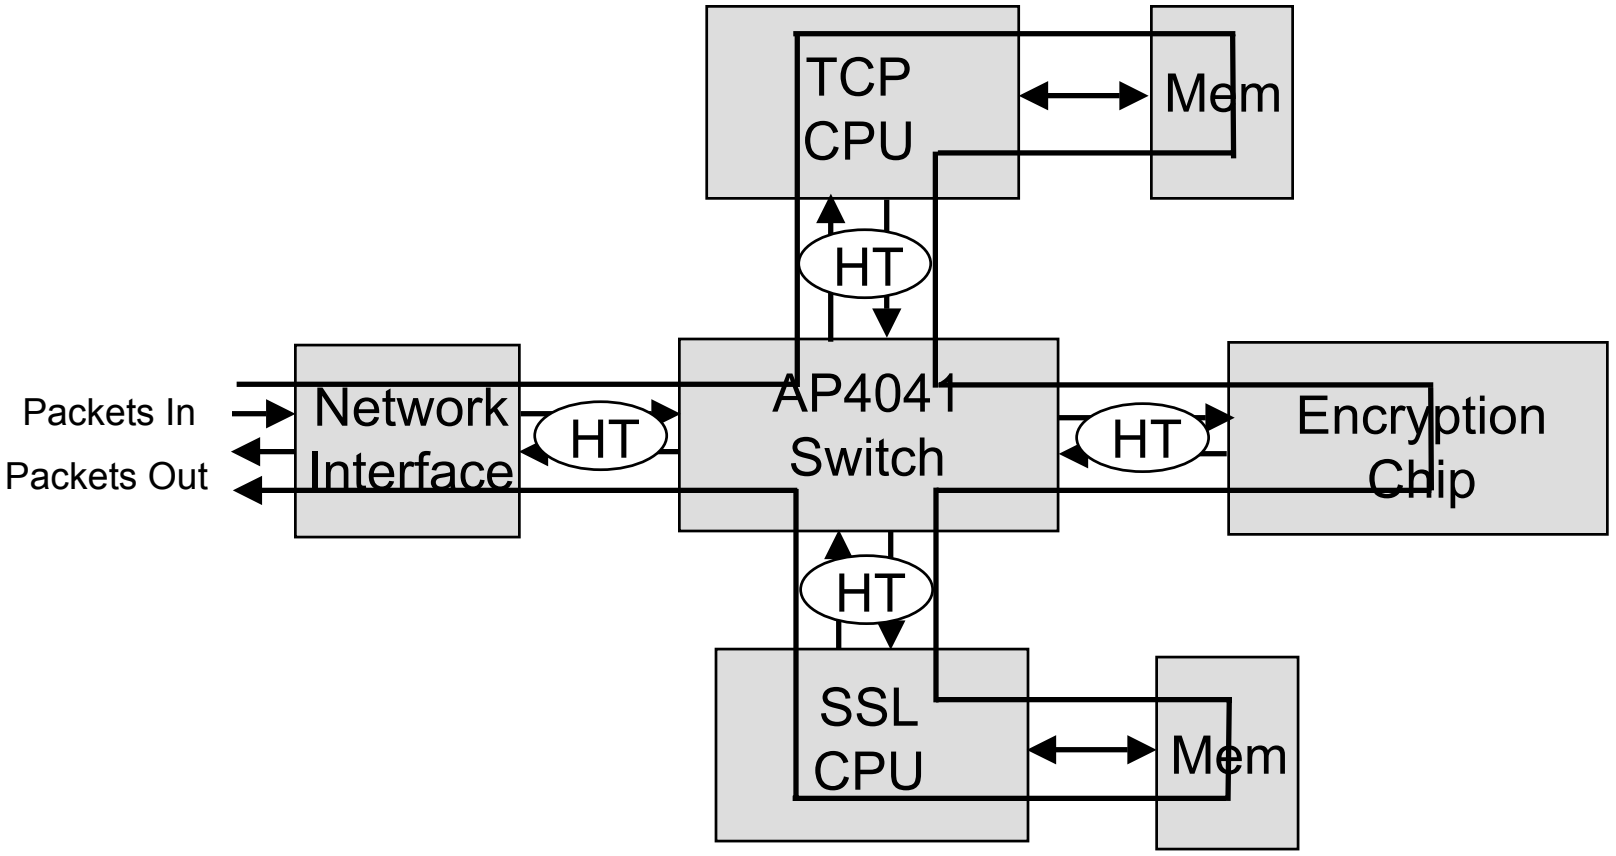

Example Networking Application: QuickLogic OC-48 WAN-to-LAN Switching

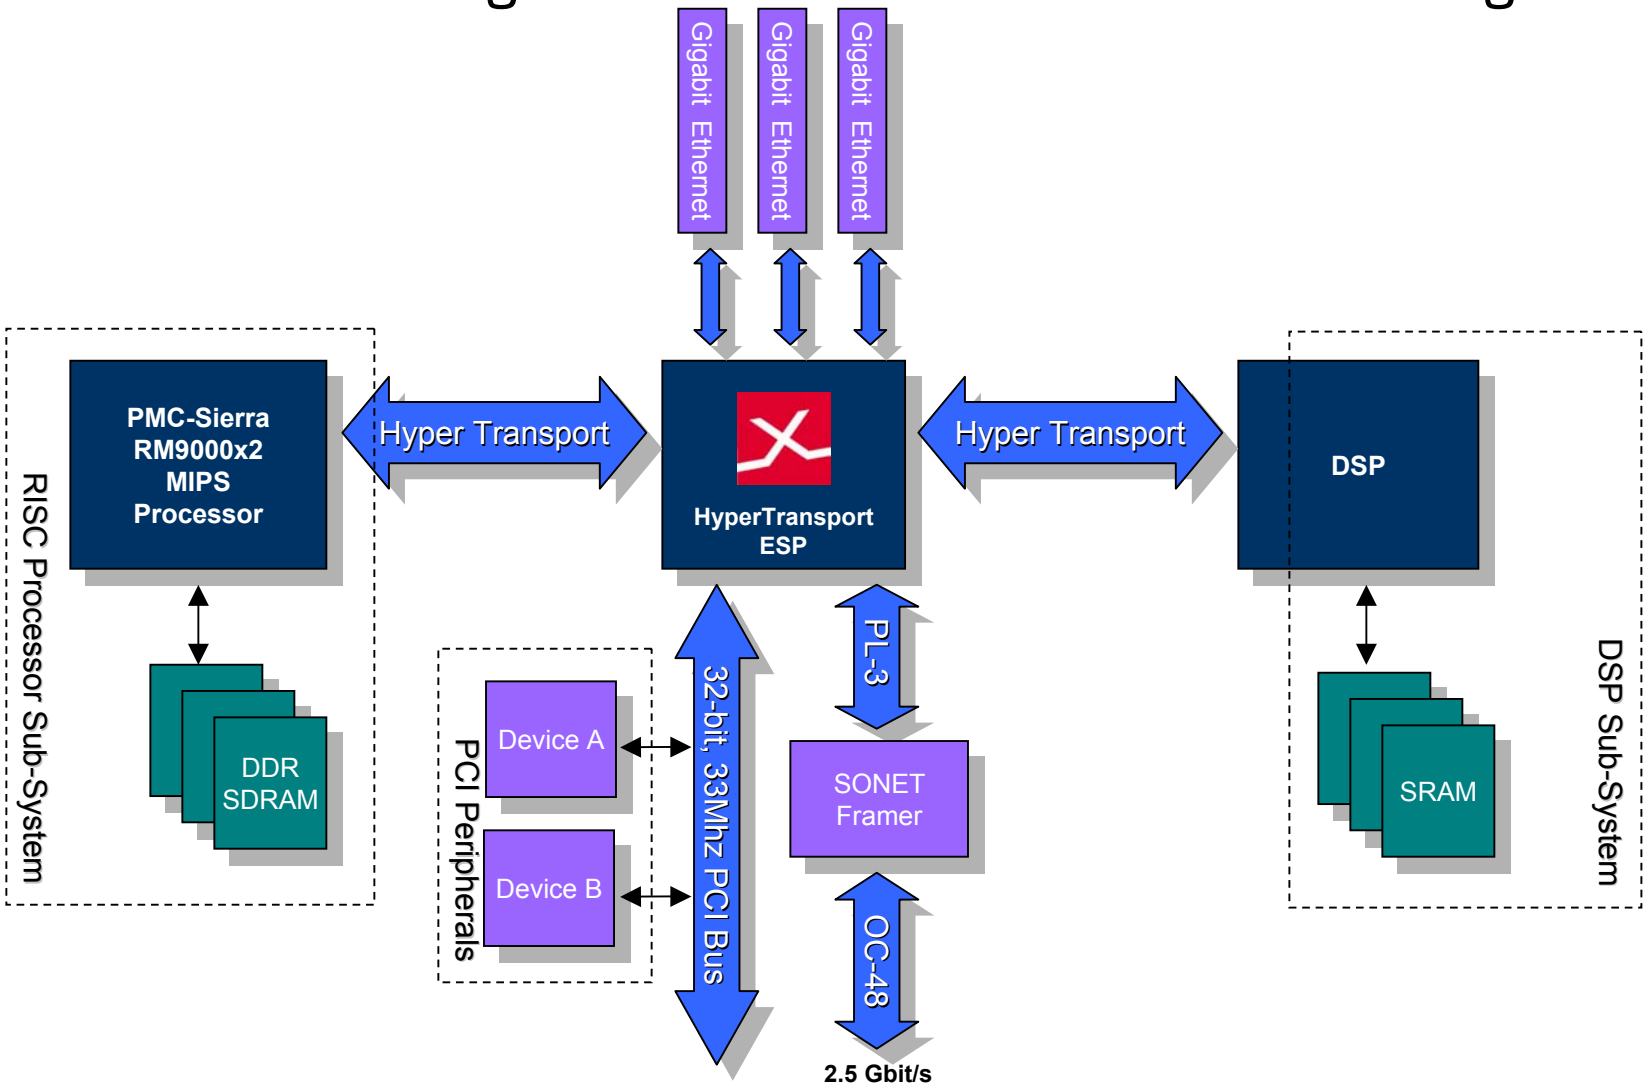

HyperTransport™ Status

- Established, mature technology
  - Broad industry support
    - Custom, ASIC and FPGA vendors
  - Some vendors are sampling HyperTransport™ devices now
  - PMC-Sierra plans to sample HyperTransport™ RM9000x2™ next year

HyperTransport™ Technology Consortium

- Independent consortium which controls HyperTransport™ specification
- Royalty-free licensing
- Open HT Specification
  - www.hypertransport.org

Concluding Thoughts

- HyperTransport™ scalability
  - Higher speeds, wider I/O
  - High performance "entry point"
  - Switching
- Rich architectural possibilities for networking applications
- Leverage-standardization across multiple industry segments
- Thriving collection of HT suppliers-IP, CPU, bridge, peripheral, ASIC, programmables.....

Contact Information

HyperTransport Consortium

- www.hypertransport.org

PMC-Sierra

- Van_Lewing@pmc-sierra.com
  or
- Info@pmc-sierra.com